

**Technical field**

The present invention pertains to a filter for an open system interface layer2 traffic separation in at least one router in a network, and a method therefore.

**Background art**

5 When deploying network devices such as routers and switches for an Ethernet ® based network or the like network, the current OSI layer2 (Open Systems Interconnection), deploying MAC addressing (Media Access Control addressing), technology enables VLANs (Virtual Local Area Networks) to be used for separating physical ports in a device, such as router and switch on layer2, and to bind ports belonging to the same VLAN together over  
10 multiple devices, so called "trunking".

On OSI layer3, deploying IP addressing through routers, each VLAN require a different IP subnet for addressing. Over the past few years several attempts have been made in using this technology to deploy a broadband network.

15 Ethernet ® is a shared media according to CSMA/CD (Carrier Sense Multiple Access with Collision Detect), which means that all hosts that are connected to one and the same Ethernet ® get all traffic, but they select it in dependence of their MAC address.

A typical broadband network consist of a number of switches or routers deployed in a residential area to connect individual households to a common infrastructure, the so called service provider infrastructure.

20 By using Ethernet technology to accomplish this, it immediately introduces a security problem of connecting different premises such as households and the like to a single shared infrastructure as Ethernet provides.

A service provider has to consider:

- 25
- Connecting each customer to a separate VLAN – thereby requiring numerous small IP subnets, one for each VLAN to preserve layer2 separation,
  - Connecting customers to a single VLAN – thereby requiring a single, larger IP subnet, but introducing the risk of allowing layer2 access between different customers, for example, Microsoft ® file-sharing.

30

To solve this filtering problem some implementations use port protection features where traffic between two ports in the same device are comprised in the same VLAN is prevented. This means that the hosts connected on those ports are unable to exchange any traffic. Further enhancements to this type of solution has included forwarding packets

between the protected ports to an upstream filtering device that makes a decision if data packet traffic should be permitted, and if so, forwarding the traffic back to its destination. This will of course put more load on the backbone link used between the switch and the filtering device.

5        With a current increase in the number of connected computers to Ethernet ® networks, problem regarding data traffic collision growths. In order to solve this problem, bridges where invented, which divide an Ethernet ® in several segments and remembered/learned in which segments the different MAC addresses resided. Thereafter forwarding of packets is only accomplished of packets that where aimed to the broadcast 10 address or to a MAC address that resides in another segment than it was transmitted from. But the different segments are still part of the same broadcast domain.

Current switches are further developments of the bridge. They could be said to have a bridge in every port. The switch remembers/learns which MAC addresses that reside on every port, respectively, and achieves forwarding between ports only if the traffic is 15 intended for a MAC address on a different port. Every port thus becomes a segment, but every port (all segments) are still a part of the same broadcast domain, as a broadcast is transmitted to every port. An advantage with a switch is that it communicates in high speeds which accomplishes that a number of ports can communicate with each other at the same time with maximum speed.

20        Switching technique has progressed, e.g. through the introduction of VLAN, trunking and spanning-tree.

VLAN makes it possible to group ports in a switch to different broadcast domains. It involves that the ports comprised in a specific VLAN are unable to communicate with ports in a different VLAN. At least not through layer2, which calls for a router to 25 connect such ports.

In RFC1027 (Request For Comment document under the control of IETF; Internet Engineering Task Force) a technique known as “Proxy-ARP” is described, in which a routing device responds to ARP requests for any address outside the local subnet requested by a locally connected host, thereby making the host send all traffic to the router without 30 requiring the understanding of an IP default-route. This was used in the early days of the Internet to guide hosts in lack of a complete understanding of IP to communicate using the IP protocol. It is rarely used today.

#### Summary of the described invention

The present invention aims to solve problems related to OSI layer2 broadcasting and the limited possibility to divide IP addresses into subsets for a plurality of VLANs.

In order to achieve its goals and aims, the present invention sets forth a filter for an open system interconnection layer2 traffic separation in at least one Access Switching

- 5 Router in a network. The ports in the routers are configured to the same virtual local area network. The filter is filtering data packet traffic to the ports. It further comprises:

means for intercepting layer2 traffic from a network connected source device for a MAC-address belonging to the virtual local area network, determining if traffic is permitted to be forwarded to other ports;

- 10 means for intercepting Address Resolution Protocol broadcasts in such traffic, responding to the broadcast to the source device regardless of if a destination device layer2 domain is the same as source device layer2 domain, the source device thus determining that the broadcast has acknowledged the layer2 address of a sought destination device, whereby the source device transmits data packets to the destination device, the router receiving the 15 transmitted data packets;

means for determining the egress port to the destination device;

means for determining the layer2 address of the destination device;

- means for adjusting the layer2 header from the received data packet, the means for setting the source layer2 address, setting the a router source address for the data 20 packets, the means for determining the layer2 address of the destination device, setting the destination layer2 address to that of the destination device, transmitting the data packet to the destination device; and

- 25 thus simulating that if the source device and destination device is in the same layer2 domain, the router layer2 address is the actual destination address both for the source and destination device, or simulating that if the source device and destination device are not in the same layer2 domain but in the same layer3 subnet, the router layer2 address is the actual destination layer2 address for the source to the destination.

- 30 In one embodiment of the present invention it is provided that a port that resides in a sub router is provided with said routers layer2 address when addressing the destination device.

Another embodiment provides that a router is investigating the source and/or destination address to determine the best exit port for the packet, to determine if the packet is in profile for rate-limiting, or to do other filtering based on information in the open system interconnection layer3 and higher protocol layers.

A further embodiment provides that the Access Switching Router is a combination of a layer2 switch and a layer3 router, combining the capabilities of layer2 switching with advanced packet control and forwarding decisions in a layer3 router.

5 A still further embodiment is providing the use of IP subnet, spreading it over several premises and a multiple of Access Switching Router and the same subnet in multiple layer2 domains, thus covering more customers. Yet another embodiment is providing a customer having multiple computers to receive more addresses.

The present invention also sets forth a method for a filter in an open system interconnection layer2 traffic separation in at least one Access Switching Router in a network.

10 10 A router having ports in the routers configured to the same virtual local area network. The filter is filtering data packet traffic to the ports. It further comprises the steps of:

intercepting layer2 traffic from a network connected source device (HostA, HostB) for a Media Access Control address belonging to the virtual local area network, determining if traffic is permitted to be forwarded to other ports;

15 15 intercepting Address Resolution Protocol broadcasts in such traffic, responding to the broadcast to the source device regardless of if a destination device layer2 domain is the same as source device layer2 domain, the source device thus determining that the broadcast has acknowledged the layer2 address of a sought destination device, whereby the source device transmits data packets to the destination device, a router receiving the transmitted data 20 packets;

determining the egress port to the destination device;

determining the layer2 address of the destination device;

adjusting the layer2 header from the received data packet, the

means for setting the source layer2 address, setting the routers source address for the data 25 packets, the means for determining the layer2 address of the destination device, setting the destination layer2 address to that of the destination device, transmitting the data packet to the destination device; and

30 30 thus simulating that if the source device and destination device is in the same layer2 domain, the router layer2 address is the actual destination address both for the source and destination device, or simulating that if the source device and destination device are not in the same layer2 domain but in the same layer3 subnet, the router layer2 address is the actual destination layer2 address for the source to the destination.

It is appreciated that the method is able to perform the steps of the attached set of dependent method claims conforming to the above described embodiments.

### Brief description of the drawings

Henceforth reference is had to the accompanying drawings for a better understanding of given examples and embodiments of the present invention, whereby:

5      Fig. 1 schematically illustrates a residential area connected to a broadband network in accordance with prior art;

Fig. 2 schematically illustrates a gateway connected between two broadband networks in accordance with prior art; and

Fig. 3 schematically illustrates a broadband network in accordance with the present invention.

### 10                  Detailed description of preferred embodiments

In order to be able to understand the solution, in accordance with the present invention, to problems related to layer2 data traffic, it is also important to understand the fundamental features of IP addressing. A fundamental part of using Ethernet ® for IP communication is the use of the ARP (Address Resolution Protocol) protocol. ARP is used to 15 resolve between OSI layer2 and layer3 addresses. It enables hosts to determine the layer2 address of another device when the layer3 address is already known. This is used when a host on an IP subnet intends to communicate with another host on that same subnet. The ARP is thus used for interpretation between layer2 addresses (Ethernet ® MAC addresses) and layer3 addresses (IP)

20                  A fundamental part of IP is that not every device in a network needs to know about a provided global routing table. If a device has a packet to forward to an unknown destination, the device may be configured with a default-route a path to use for any traffic for which there is not an explicit route. The default route is always an IP address on a subnet that the host is directly attached to. The layer2 address of the default route is remembered/learned 25 by the ARP protocol unless it is not statically configured in the host.

In accordance with the present invention, a router is defined as a device that analyses OSI layer 3 or higher protocol information to make a traffic forwarding decision.

This includes but is not limited to investigating the source and/or destination address to determine the best exit port for the packet, to determine if the packet is in profile 30 for rate-limiting, or to do other filtering based on information in the OSI layer3 and higher protocol layers.

The Access Switching Router (ASR) is a combination of a layer2 switch and a layer3 router. It combines the capabilities of layer2 switching with advanced packet control

and forwarding decisions in a layer3 router. This definition fits the definition of a router in accordance with the present invention and also incorporates the unique filtering features described herein.

The advantages of the present invention enables all Ethernet ® ports on the ASR  
5 to be configured to the same VLAN, which enables the ports to share the same IP subnet. Hence, no dividing of the subnet, for example, a 32 bit IP address, has to take place. Every time a subnet is created two addresses disappear. Those are the so called net address and the address being the subnets broadcasting address. When corporations, Internet service providers etc. connect to Internet they apply for IP addresses. An assignment of addresses is dependent  
10 on how many computers that are connected to the network, how the network is to be designed and its pace of growth in the following years.

Given an example, a company is assigned 192.168.1.0/24 as address, where /24 denotes the dimension of the subnet. As IP addresses have 32 binary bits, it is easier to provide an example in binary notation:

15

192.168.1.0 = 11000000 10101000 00000001 00000000  
/24 equals a one decimal subnet-mask of 255.255.255.0 binary reassembling  
11111111 11111111 11111111 00000000

20

The part of a subnet where the subnet-mask is 0, below denoted the host part, is the part that is allowed to use for setting an IP address for the single computers. The part where the subnet-mask is 1 must always be the same. Two addresses in this part may never be used for computers and these are the net-number itself when the host part only comprises binary 0, and the broadcast address when the host part only comprises binary 1. Hence:

25

11000000 10101000 00000001 00000000 192.168.1.0  
11000000 10101000 00000001 11111111 192.168.1.255

It is not likely that 250 computers are connected to one and the same segment.  
30 Probably it consists of several segments divided into several layer2 broadcast domains, thus every layer2 domain needs one IP subnet of its own. Therefore it is necessary to divide the 256 addresses in smaller subnets. This is accomplished by further prolonging the subnet-mask, i.e., the part comprising binary 1.

Exampel:

11000000 10101000 00000001 00000000 192.168.1.0  
11111111 11111111 11111111 11000000 255.255.255.192

5

The subnet-mask is now intruding on two bits in the last octet. This means that there are 6 bits left for a host address which decimaly reassembles 64. Hence, the 256 addresses have turned into four subnets of each 64 addresses.

10 11000000 10101000 00000001 00000000 192.168.1.0  
11111111 11111111 11111111 11000000 255.255.255.192

11000000 10101000 00000001 01000000 192.168.1.64  
11111111 11111111 11111111 11000000 255.255.255.192

15 11000000 10101000 00000001 10000000 192.168.1.128  
11111111 11111111 11111111 11000000 255.255.255.192

11000000 10101000 00000001 11000000 192.168.1.192  
20 11111111 11111111 11111111 11000000 255.255.255.192

Each and every one of these four subnets are having two addresses that are not allowed to use. Decimaly, they are:

25 Subnet 192.168.1.0 forbidden 192.168.1.0 and 192.168.1.63  
Subnet 192.168.1.64 forbidden 192.168.1.64 and 192.168.1.127  
Subnet 192.168.1.128 forbidden 192.168.1.128 and 192.168.1.191  
Subnet 192.168.1.192 forbidden 192.168.1.192 and 192.168.1.255

30 Binary reassembling:

11000000 10101000 00000001 00000000 192.168.1.0  
11111111 11111111 11111111 11000000 255.255.255.192  
11000000 10101000 00000001 00111111 192.168.1.63

11111111 11111111 11111111 11000000 255.255.255.192

11000000 10101000 00000001 01000000 192.168.1.64

11111111 11111111 11111111 11000000 255.255.255.192

5 11000000 10101000 00000001 01111111 192.168.1.127

11111111 11111111 11111111 11000000 255.255.255.192

11000000 10101000 00000001 10000000 192.168.1.128

11111111 11111111 11111111 11000000 255.255.255.192

10 11000000 10101000 00000001 10111111 192.168.1.191

11111111 11111111 11111111 11000000 255.255.255.192

11000000 10101000 00000001 11000000 192.168.1.192

11111111 11111111 11111111 11000000 255.255.255.192

15 11000000 10101000 00000001 11111111 192.168.1.255

11111111 11111111 11111111 11000000 255.255.255.192

It is now possible to divide one of these 64 address subnets in two parts, receiving two subnets of 32 addresses, but which each comprise two forbidden addresses:

20

11000000 10101000 00000001 11000000 192.168.1.192

11111111 11111111 11111111 11100000 255.255.255.224

11000000 10101000 00000001 11011111 192.168.1.223

11111111 11111111 11111111 11100000 255.255.255.224

25

11000000 10101000 00000001 11100000 192.168.1.224

11111111 11111111 11111111 11100000 255.255.255.224

11000000 10101000 00000001 11111111 192.168.1.255

11111111 11111111 11111111 11100000 255.255.255.224

30

In a broadband network 32 addresses are in excess for a single household. Every computer connected to a subnet is deemed to have an address, which also includes the default-gateway router, there is a demand of at least two addresses for every household, one for the computer and one for the router. If the household is in control of more than one computer, a bigger subnet is needed.

Therefore, two addresses per household requires that the smallest subnet has to have the dimension of four addresses. Binary:

5    11000000 10101000 00000001 10000000 192.168.1.0  
      11111111 11111111 11111111 11111100 255.255.255.252

Since two addresses are forbidden:

10    11000000 10101000 00000001 00000000 192.168.1.0  
      11111111 11111111 11111111 11111100 255.255.255.252

11000000 10101000 00000001 00000011 192.168.1.3  
      11111111 11111111 11111111 11111100 255.255.255.252

15    the addresses left to use are 192.168.1.1 and 192.168.1.2. In the next subnet, the addresses 192.168.1.4 and 192.168.1.7 are forbidden. Addresses that can be used are 192.168.1.5 and 192.168.1.6 and so forth.

Out of the 256 addresses from the start there are  $256/4 = 64$  subnets or 64 customers. One half of the addresses in these kind of small subnets are retained as broadcast and net addresses, and the loss of address space is 50%.

If subnets are designed in bigger dimensions, the loss of address space decreases due to broadcast and net addresses (8 addresses per subnet provides  $256/8 = 32$  subnets, a 25 % loss of address space). But there are 6 useful addresses per subnet, and if the router is provided one, there are 5 addresses per household. If those 5 addresses are not fully used, because there are not more than two computers in every household, there still is an address loss as 3 addresses are not used.

Through the solution in accordance with one embodiment of the present invention, it is enabled to use 254 addresses of the 256 provided in the subnet and spread it over several premises and multiple ASRs thus covering more customers. If one customer has more computers than another customer, no extra loss of address space is introduced as the customer with the greater number of computers receives more addresses. Therefore, the loss of address space with the present invention is held at a few percentages if the network is built to optimize the address space.

According to the present invention a filter is applied which hinders any layer2 traffic between the ports belonging to the VLAN, except traffic with protocol options indicating that the data carried in the layer2 packet is IP, IPv6 or any other traffic acceptable for the purpose of communication. This means that even though the ports belong to the same 5 layer2 broadcast domain, traffic between them is prevented from being switched based on their source and destination layer2 address.

When a client attached to a port starts to transmit, the first packet will traverse the Ethernet ® segment, including the ASR.

Whenever the client host seeks to communicate with another host it will issue an 10 ARP request for either the default-route, if the destination is not part of the client hosts IP subnet, or the destination itself, if its destination address is on the client hosts same subnet. This ARP request is a layer2 broadcast, which typically traverses the entire VLAN. The ARP message is intercepted, in accordance with the present invention, by the ASR and prevented from being forwarded to any other port belonging to that VLAN. If the ARP request is for a 15 destination that is present on any other port on the ASR or if the destination is known in the ASR layer3 routing table, the ASR is responding to the ARP request with its own MAC-address as next-hop. This procedure makes the client host believe, simulates, that the ASR layer2 address is the destination layer2 address to be used to reach the real layer3 destination. Thus, the client host transmits the packet to the ASR layer2 address.

20 If the packet is determined to be forwarded out on another of the ASR ports, based on the destination layer3 address and the content of the ASR routing table and/or address resolution table, the source-MAC address of the packet is changed to the ASR layer2 address on the egress port. The source IP address will continue to be that of the original client host address. Thereby, the receiver in the ASR remembers/learns that the source client host 25 address maps to the ASR layer2 address and any return traffic to the source client host is directed to the ASR rather than directly to the source client MAC address. In this manner both the source and the destination client hosts are simulated to believe that the ASR MAC-address is the address of the other host and communication flow is maintained.

To be able to communicate with TCP/IP a host has to be configured with:

30

- an IP address
- a subnet-mask
- a default-gateway
- a name server

A name server is used to connect between names and IP addresses on the Internet.

Fig. 1 schematically illustrates a residential area connected to a broadband network 10 in accordance with prior art. At switch 12 is depicted a VLAN with all ports 14 connected to it, meaning that neighbours have layer2 access between themselves. This enables one neighbour to for example browse another neighbours hard-drive. The switch 16 comprises that every port 14 belongs to a different VLAN, which requires a small IP subnet per VLAN. This is a waste of address space because every subnet introduces unusable addresses for the network and the broadcast feature. A subnet with two usable addresses also requires two unusable addresses, wasting 50 % of the address space. The devices 18 in Fig. 1 are routers.

Fig. 2 schematically illustrates a gateway 30 connected between two broadband networks 32, 34 in accordance with prior art, also depicting HostA, HostB and HostC.

The following sequence describes the conventional operation of the ARP routing protocol.

The first sequence of steps 1)-9) provides an example where HostA transmits to HostB with reference to Fig. 2:

- 1) HostA has IP packet to send
- 2) HostA compares HostAs address + subnetmask with HostBs address
- 20       3) HostB is on same network as HostA
- 4) HostA sends ARP broadcast to Network1 requesting HostBs layer2 address.
- 5) HostB recognize request for its layer2 address
- 6) HostB responds
- 25       7) HostA now has HostBs layer2 address
- 8) HostA transmit data
- 9) HostB receives data

The second sequence of steps 1)-17) provides an example where HostA transmits to HostC with reference to Fig. 2:

- 1) HostA has IP packet to send
- 2) HostA compares HostAs address + subnetmask with HostCs address
- 3) HostC is not on same network as HostA

- 4) HostA sends ARP broadcast to Network1 requesting Gateways layer2 address
- 5) Gateway recognize request for its layer2 address
- 6) Gateway responds
- 5 HostA now has Gateways layer2 address
- 8) HostA transmit data
- 9) Gateway receive data
- 10) Gateway strips away layer2 information from packet
- 11) Gateway looks up HostC address in routing table and determines egress interface
- 10 12) Gateway send ARP broadcast to Network2 requesting HostCs layer2 address
- 13) HostC recognize request for its layer2 address
- 14) HostC responds
- 15) Gateway now has HostCs layer2 address
- 15 16) Gateway builds new layer2 header for packet and transmit data
- 17) HostC receives data.

If the gateway 30 had not been directly connected to Network2, step12 would  
20 instead have been “forwarding the packet towards Network2”, repeating steps 9,10, 11 and the new step 12 in every gateway along the path until the gateway that is connecting directly to Network2, receiving the packet where steps 12-17 according to the flow above would commence.

Fig. 3 schematically illustrates a broadband network 40 in accordance with the  
25 present invention, having two ASR routers 42, 44. HostA and HostB are connected to router 42 and HostC connected to router 44. Both routers 42 and 44 have a direct connection between each other, where router 42 comprises the filter of the present invention. Fig. 3 also depicts a HostD connected to the broadband network via Internet.

The filter of the present invention is provided for an open system  
30 interconnection layer2 traffic separation in at least one ASR router 42 in a broadband network 40. All ports (not shown) in the routers 42, 44 are configured to the same VLAN. ASR 44 is a sub router to router 42 or just connected and provides the same filtering advantages in accordance with the present invention. Data packet traffic is intercepted by the router 42 comprising the filter, which is filtering data packet traffic to the ports. The filter comprises:

means for intercepting layer2 traffic from a network connected source device (HostA, HostB) for a MAC-address belonging to the virtual local area network, determining if traffic is permitted to be forwarded to other ports;

- means for intercepting Address Resolution Protocol broadcasts in such traffic,
- 5 responding to the broadcast to the source device regardless of if a destination device layer2 domain is the same as source device layer2 domain, the source device thus determining that the broadcast has acknowledged the layer2 address of a sought destination device, whereby the source device transmits data packets to the destination device, the router receiving the transmitted data packets;
- 10 means for determining the egress port to the destination device;
- means for determining the layer2 address of the destination device;
- means for adjusting the layer2 header from the received data packet, the means for setting the source layer2 address, setting the routers source address for the data packets, the means for determining the layer2 address of the destination device, setting the 15 destination layer2 address to that of the destination device, transmitting the data packet to the destination device.

The filter of the present invention is thus simulating that if the source device and destination device is in the same layer2 domain, the router layer2 address is the actual destination address both for the source and destination device, or simulating that if the source 20 device and destination device are not in the same layer2 domain but in the same layer3 subnet, the router layer2 address is the actual destination layer2 address for the source to the destination.

It is appreciated that the means of the present invention preferably are software building blocks in a router or a combination of hardware and software.

25 In the following three scenarios for packet flow in accordance with the present invention and with reference to Fig. 3 are provided.

It is to be noted that in IP routing, the encapsulation and decapsulation of layer2 headers on an IP packet is a conventional procedure. The IP header with the IP source and destination address is left untouched while the layer2 headers for Ethernet, TokenRing, 30 FrameRelay, ATM or other layer2 technology that is used changes. Because the layer2 protocol is not routable, the source address is always set to that of the device transmitting the packet. This is conventional.

The first scenario with sequence steps 1) to 13) describes packet transmission from HostA to HostB. Both hosts are connected to ports in the same ASR. The ports are

configured to belong to the same broadcast domain (VLAN) but port protection with additional features is enabled on the ASR in accordance with the present invention.

First scenario

- 5            1) HostA has IP packet to send
- 2) HostA compares its address + subnetmask with HostB and determines they are on the same subnet.
- 3) HostA sends ARP broadcast for HostBs address
- 4) Because of filters between the ASR 42 ports, the broadcast cannot reach HostB.
- 10            5) The ASR intercepts the ARP broadcast and determines it knows where HostB is located.
- 6) The ASR responds to the ARP request for HostB, setting its own layer2 address as the address for HostB
- 15            7) HostA receives the ARP response and think it now know the layer2 address for HostB.
- 8) HostA transmit data
- 9) ASR 42 receive data.
- 10) ASR 42 removes layer2 information and determine the egress port for HostB
- 20            11) ASR 42 sets its own layer2 address as source for the packet and encapsulates the packet for HostB.
- 12) ASR 42 transmit data
- 13) HostB receives the data from HostA.

25

Because that the ASR 42 layer2 address is set as source, HostB believes that the layer2 address of ASR 42 is that of HostA. Likewise, due to the ARP response, HostA will believe that the layer2 address of ASR 42 is that of HostB.

The second scenario with sequence steps 1) to 18) describes packet transmission from HostA to HostC. The hosts are connected to ports on different ASRs. But the address sharing features of the ASR and central management system agreed the hosts to receive IP addresses by DHCP from the same IP subnet. The ASRs have exchanged routing information informing each other about connected hosts.

Second scenario

- 1) HostA has IP packet to send
- 2) HostA compares its address + subnetmask with HostC and determines they are on the same subnet.
- 5 3) HostA sends ARP broadcast for HostCs address
- 4) Because of filters between ASR 42 ports, the broadcast do not reach any other port on the ASR.
- 5) ASR 42 intercepts the ARP broadcast and determines it knows where HostC is located.
- 10 6) ASR 42 responds to the ARP request for HostC, setting its own layer2 address as the address for HostC
- 7) HostA receives the ARP response and think it now know the layer2 address for HostC.
- 8) HostA transmits the packet
- 15 9) ASR 42 receives the packet.
- 10) ASR 42 removes layer2 information and determine the egress port for HostC.
- 11) ASR 42 encapsulates the packet with appropriate layer2 headers for the link to ASR 44.
- 20 12) ASR 42 forwards the packet towards ASR 44
- 13) ASR 44 receives the packet.
- 14) ASR 44 removes layer2 encapsulation used on the link from ASR 42.
- 15) ASR 44 determines the egress port for the packet towards HostC.
- 16) ASR 44 encapsulates the packet with layer2 headers, setting its own layer2 address as source.
- 25 17) ASR 44 transmit data
- 18) HostC receives the data from HostA.

Because of that the ASR 42 is responding to the ARP request, HostA will believe that the layer2 address of ASR 42 is that of the HostC. Because of the ASR 44 setting its layer2 address as source for the packet to HostC in the final steps above. HostC thus believes that the layer2 address of ASR 44 is that of the HostA.

The third scenario with sequence steps 1) to 15) describes packet transmission from HostA to HostD. HostA is connected to a port on ASR 42. HostD is connected somewhere on the Internet.

5

Third scenario

- 1) HostA has IP packet to send
- 2) HostA compares its address + subnetmask with HostD and determines they are not on the same subnet.
- 3) HostA sends ARP broadcast for default-gateway address
- 10 4) Because of filters between the ASR 42 ports, the broadcast cannot reach any other port on the ASR.
- 5) The ASR intercepts the ARP broadcast and determines it is the default-gateway.
- 15 6) The ASR responds to the ARP request for default-gateway with its own layer2 address.
- 7) HostA receives the ARP response and think it now know the layer2 address for the defalt gateway
- 8) HostA transmit data
- 9) ASR 42 receive data.
- 20 10) ASR 42 removes layer2 information and determine the egress port for HostD.
- 11) ASR 42 encapsulates the packet with appropriate layer2 headers for the link towards HostD
- 12) Gateways along the path between ASR 42 and HostD repeat steps 9-11.
- 25 13) The gateway connecting HostD receives the packet
- 14) The gateway performs ARP lookup and forwards the packet towards HostD according to Internet standards.
- 15) HostD receives the data.

30

The present invention has been described through examples and embodiments not intended to limit the scope of protection, whereby a person skilled in the art is able to derive further embodiments by the attached set of claims.